arXiv:1510.01495v1 [cs.IT] 6 Oct 2015

Quantifying Emergent Behavior of Autonomous Robots

Georg Martius[1]* and Eckehard Olbrich[2]

[1] IST Austria, Am Campus 1, 3400 Klosterneuburg, Austria

[2] Max Planck Institute for Mathematics in the Sciences, Inselstr. 22, 04103 Leipzig, Germany

October 7, 2015

Abstract

Quantifying behaviors of robots which were generated autonomously from task-independent objective functions is an important prerequisite for objective comparisons of algorithms and movements of animals. The temporal sequence of such a behavior can be considered as a time series and hence complexity measures developed for time series are natural candidates for its quantification. The predictive information and the excess entropy are such complexity measures. They measure the amount of information the past contains about the future and thus quantify the nonrandom structure in the temporal sequence. However, when using these measures for systems with continuous states one has to deal with the fact that their values will depend on the resolution with which the systems states are observed. For deterministic systems both measures will diverge with increasing resolution. We therefore propose a new decomposition of the excess entropy in resolution dependent and resolution independent parts and discuss how they depend on the dimensionality of the dynamics, correlations and the noise level. For the practical estimation we propose to use estimates based on the correlation integral instead of the direct estimation of the mutual information using the algorithm by Kraskov et al. (2004) which is based on next neighbor statistics because the latter allows less control of the scale dependencies. Using our algorithm we are able to show how autonomous learning generates behavior of increasing complexity with increasing learning duration.

1 Introduction

Whenever we create behavior in autonomous robots we strive for a suitable measure in order to quantify success, learning progress or compare algorithms with each other. When a specific task is given, then typically the task provides a natural measure of success, for instance walking may be measured by velocity and perturbation stability. In cases where behavior is learned via optimization of a global objective function then the same function can also be used as a quantification, so creating and quantifying behavior often go hand in hand. This also applies in principle to behavior from task-independent[1] objectives that have been recently more and more successful in generating emergent autonomous behavior in robots [1, 14, 24, 32, 35, 39, 48]. However, there are several cases of emergent behavior where this strategy fails: if the behavior arises from a) optimizing a local function [7, 28], b) optimizing a crude approximation of a computationally expensive objective function [28], c) local interaction rules without an explicit optimization function [8, 9, 37], and d) a biological system (e. g. freely moving animals) where we don't know the underlying optimization function. Thus, independent of the origin of behavior it would be useful to have a quantitative description of its structure. This would allow to objectively compare algorithms with each other and to compare technical with biological autonomous systems. This paper presents a measure of behavioral complexity that is suitable for the analysis of this kind of emergent behavior.

We base our measure on the predictive information (PI) [2] of a set of observables, such as joint angles. The PI is the mutual information between past and future of a time-series and captures how much information (in bits) can be used from the past to predict the future. It is also closely linked to the excess entropy [5, 41] or effective measure complexity [18], and it is a natural measure for the complexity of a time series because it provides a lower bound for the information necessary for optimal prediction. Thus it is a promising choice as a measure. Given these favorable properties, the PI was also proposed as an intrinsic drive for behavior generation in autonomous robots [1, 28, 48]. On an intuitive level maximizing the PI of the sensor process, leads to a high variety of sensory input — large marginal

[1] A task is here understood as a specific desired target behavior such as walking or reaching.

entropy — while keeping a temporal structure — high predictability corresponding to a small entropy rate.

Unfortunately, there are conceptual and practical challenges in using the PI for generating and measuring autonomous behavior. Conceptually, for systems with continuous values the PI provides a family of functions depending on the resolution of the measurements and practically it is difficult to estimate the PI. For a fixed partition (single resolution) in a low-dimensional case it is possible to estimate the one-step PI and adapt a controller to maximize it [48]. In high-dimensional systems it cannot be used directly in an online algorithm. An alternative approach [28] uses a dynamical systems model, linearization and a time-local version of the PI [28] to obtain an explicit gradient for locally maximizing PI for continuous and high-dimensional state spaces (no global optimization). A simplified version of the resulting neural controller [8] was used for generating the examples presented in this paper.

The aim of the present paper is to provide a measure of behavioral complexity that can be applied to emergent autonomous behavior, animal behavior and potentially other time-series. In order to do this it turns out that the predictive information as a complexity measure has to be refined by taking into account its scale dependency. Because the systems of interest are high-dimensional we investigate different methods to estimate information theoretic quantities in a resolution dependent way and present results for different robotic experiments. The emergent behavior in these experiments are generated by the above mentioned learning rules. On the one hand we want to increase our understanding of the learning process guided by the time-local version of the PI and on the other hand demonstrate how the resolution dependence allows us to quantify behavior on different length-scales.

We argue that many natural behaviors are low-dimensional, at least on a coarse scale. For instance in walking all joint angles can be typically related to a single phase, so in the extreme case it is one-dimensional [45]. However, as we show below in 2.1, this contradicts to the global maximization of the PI which would maximize the dimension. We discuss this contradiction and how it was circumvented in the practical applications. Along with the paper there is a supplementary material located at `playfulmachines.com/QuantBeh2015` and the source code can be found in the repository `github.com/georgmartius/behavior-quant`.

The plan of the paper is as follows. In Sec. 2.1 we introduce the information theoretic quantities and in Sec. 2.3 the two estimation methods used in the paper. In Sec. 2.4 we describe the algorithm used to calculate the quantification measures. In Sec. 2.5 we demonstrate their behavior at the example of the Lorenz system including the effects of noise. Sec. 3 starts with explaining the robots and the control framework used for the experiments which are described in 3.2. The results of the experiments quantification are presented in Sec. 3.3 and are discussed in Sec. 4.

2 Methods

2.1 Entropy, dimension and excess entropy

Starting point is a (in general) vector valued stationary time series $\{y_0, y_1, y_2, \ldots, y_N\}$. It is used to reconstruct the phase space of the underlying dynamical system by using delay embedding: $\vec{x}_t = \{y_t, y_{t-\tau}, y_{t-2\tau}, \ldots, y_{t-(m-1)\tau}\}$. If the original time series was $d-$dimensional the reconstructed phase space will be of dimension $d \times m$ [2]. In the following, we will only consider the case $d = 1$, i.e. scalar time series. The generalization to vector-values time series is straight forward. Let us assume that we are able to reconstruct approximately the joint probability distribution $p(\vec{x}_t, \vec{x}_{t-1}, ...)$ in a reconstructed phase space. Then we can characterize its structure using measures from information theory. Information theoretic measures represent a general and powerful tool for the characterization of the structure of joint probability distributions [19, 33]. The uncertainty about the outcome of a single measurement of the state, i.e. about $\vec{x}_t$ is given by its *entropy*. For discrete-valued random variable X with values $x \in \mathcal{X}$ and a probability distribution $p(x)$ it is defined as

$$H(X) = -\sum_{x \in \mathcal{X}} p(x) \log p(x) \ .$$

An alternative interpretation for the entropy is the average number of bits required to encode a new measurement. In our case, however, the y_t are considered as continuous-valued observables (that are measured with a finite resolution). For continuous random variables with a probability density $\rho(x)$ one

[2] There are more general embeddings possible, e.g. involving different delays or different embedding dimensions for each component of the time series.

can also define an entropy, the *differential entropy*

$$H^C(X) = -\int \rho(x) \log \rho(x) dx \ . \tag{1}$$

However, it behaves differently than its discrete counterpart: It can become negative and it will get even minus infinity if the probability measure for X is not absolutely continuous w. r. t. to the Lebesgue measure – for instance, in the case of the invariant measure of a deterministic system with an attractor dimension smaller than the phase space dimension. Therefore, when using information theoretic quantities for characterizing dynamical systems researchers often prefer using the entropy for discrete-valued random variables. In order to use them for dynamical systems with continuous variables usually either partitions of the phase space or entropy-like quantities based on coverings are employed. These methods do not explicitly reconstruct the underlying invariant measure, but exploit the neighbor statistics directly. Alternatively one could use direct reconstructions using kernel density estimators [34] or methods based on maximizing relative entropy [17, 42] to gain parametric estimates. These reconstructions, however, will always lead to probability densities, and are not suitable for reconstructing fractal measures which appear as invariant measures of deterministic systems.

In this paper we use estimators based on coverings, i. e. correlation entropies [44] and nearest neighbors based methods [26] considered in 2.3 below. For the moment let us consider a partition of the phase space into hypercubes with side-length ϵ. For a more general definition of an ϵ-partition [see 16]. In principle one might consider scaling the different dimensions of y_t differently, but for the moment we assume that the time series was measured using an appropriate rescaling. The entropy of the state vector $\vec{x}_t$ observed with a ε-partition will be denoted in the following as $H(\vec{X}; \varepsilon)$ with ϵ parameterizing the resolution.

How does the entropy change if we change ε? The uncertainty about the potential outcome of a measurement will increase if the resolution of the measurement is increased, because of the larger number of potential outcomes. If $\vec{x}$ is a m-dimensional random variable and distributed according to a corresponding probability density function $\rho(x)$ we have asymptotically for $\epsilon \to 0$ ([see Cover and Thomas 4, ch.8, p.248, theorem 8.3.1] or [16]).

$$H(X; \varepsilon) \approx H^C(X) - m \log \varepsilon \ . \tag{2}$$

This is what we would expect for a stochastic system. However, if we observe a deterministic system the behavior of an observable depends how its dimension relates to the attractor dimension. If the embedding dimension is smaller than the attractor dimension the deterministic character will not be resolved and Eq. 2 still applies. However, if the embedding dimension is sufficiently high ($m > D$ [11]) then instead of a density function $\rho(x)$ we have to deal with a D-dimensional measure $\mu(x)$ and the entropy will behave as

$$H(X; \varepsilon) \approx \text{const} - D \log \varepsilon \ . \tag{3}$$

If an behavior such as in Eqs. (2) or (3) is observed for a range of ϵ values we will call this range a *stochastic* or *deterministic scaling range*, respectively.

Conditional entropy and mutual information Let us consider two discrete-valued random variables X and Y with values $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, respectively. Then the uncertainty of a measurement of X is quantified by $H(X)$. Now we might ask, what is the average remaining uncertainty about X if we have seen already Y? This is quantified by the *conditional entropy*

$$H(X|Y) = H(X, Y) - H(Y) \ . \tag{4}$$

The reduction of uncertainty about X knowing Y is the information that Y provides about X and is called the *mutual information* between X and Y

$$I(X : Y) = H(X) - H(X|Y) \ . \tag{5}$$

Having defined the ε-dependent state entropy $H(\vec{X}; \varepsilon)$ we can now ask, how much information the present state contains about the state of the system at the next time step. The answer is given by the mutual information between $\vec{X}_t$ and $\vec{X}_{t+1}$:

$$I(\vec{X}_t : \vec{X}_{t+1}; \varepsilon) = H(\vec{X}_{t+1}; \varepsilon) - H(\vec{X}_{t+1}|\vec{X}_t; \varepsilon) \ . \tag{6}$$

Using Eq. 2 one see, that for stochastic systems the mutual information will remain finite In the limit $\varepsilon \to 0$ and can be expressed by the differential entropies:

$$I(\vec{X}_t : \vec{X}_{t+1}) := \lim_{\varepsilon \to 0} I(\vec{X}_t : \vec{X}_{t+1}; \varepsilon) = H^C(\vec{X}_{t+1}) - H^C(\vec{X}_{t+1}|\vec{X}_t) \tag{7}$$

Note, that this mutual information is invariant with respect to coordinate transformation of the system state, i.e. if $\vec{z}_j = f(x_j)$ is a continuous and invertible function, then

$$I(\vec{X}_t : \vec{X}_{t+1}) = I(\vec{Z}_t : \vec{Z}_{t+1}) \; . \tag{8}$$

However, in the case of a deterministic system, the mutual information will diverge

$$I(\vec{X}_t : \vec{X}_{t+1}; \varepsilon) \propto -D \log \varepsilon \qquad \text{for } \varepsilon \to 0 \; . \tag{9}$$

This is reasonable behavior because in principle the actual state contains an arbitrary large amount of information about the future. In practice, however, the state is known only with a finite resolution determined by the measurement device or the noise level.

Predictive information, excess entropy and entropy rate The unpredictability of a time series can be characterized by the conditional entropy of the next state given the previous states. In the following we will use an abbreviated notation for these conditional entropies and the involved entropies:

$$H_m(\vec{X}_t; \varepsilon) := H(\vec{X}_t, \vec{X}_{t-1}, \ldots, \vec{X}_{t-m+1}; \varepsilon) \tag{10}$$

$$h_m(\vec{X}_t; \varepsilon) := H(\vec{X}_t | \vec{X}_{t-1}, \ldots, \vec{X}_{t-m}; \varepsilon) \tag{11}$$

$$h_0(\vec{X}_t; \varepsilon) := H(\vec{X}_t; \varepsilon) \; .$$

The *entropy rate* ([see Cover and Thomas 4, ch 4.2]) is this conditional information if we condition on the infinite past

$$h_\infty(\vec{X}_t; \varepsilon) = \lim_{m \to \infty} h_m(\vec{X}_t; \varepsilon) \tag{12}$$

In the following we assume stationarity, i.e. we have no explicit time dependence of the joint probabilities and therefore also of the entropies. Moreover, if it is clear from the context, which stochastic process is considered, we will write $H_m(\varepsilon)$ and $h_m(\varepsilon)$ instead of $H_m(\vec{X}_t; \varepsilon)$ and $h_m(\vec{X}_t; \varepsilon)$, respectively and it holds

$$h_m(\varepsilon) = H_{m+1}(\varepsilon) - H_m(\varepsilon) \; . \tag{13}$$

For deterministic systems the entropy rate will converge in the limit $\varepsilon \to 0$ to the Kolmogorov-Sinai (KS-)entropy [3, 43] which is a dynamical invariant of the system in the sense that it is independent on the specific state space reconstruction. Moreover, already for finite $m \geq D$, $h_m(\varepsilon)$ will not depend on ε for sufficiently small ε because of (3).

To quantify the amount of predictability in a state sequence one might consider subtracting the unpredictable part from the total entropy of a state sequence. By doing this one ends up with a well known complexity measure for time series, the *excess entropy* [5, 41] or *effective measure complexity* [18]

$$E(\varepsilon) = \lim_{m \to \infty} E_m(\varepsilon) \tag{14}$$

with

$$E_m(\varepsilon) := H_m(\varepsilon) - m h_{m-1}(\varepsilon) \; . \tag{15}$$

The excess entropy provides a lower bound for the amount of information necessary for an optimal prediction. For deterministic systems, however, it will diverge because $H_m(\epsilon)$ will behave according to (3) and $h_{m-1}(\varepsilon)$ will become ϵ-independent for sufficiently large m and small ϵ, respectively

$$E_m(\varepsilon) \approx \text{const} - D \log \varepsilon \; , \tag{16}$$

with D being the attractor dimension.

The *predictive information* [2] is the mutual information between the semi-infinite past and the future time series

$$\mathrm{PI}(\varepsilon) = \lim_{m \to \infty} \mathrm{PI}_m(\varepsilon) \tag{17}$$

with

$$\begin{aligned} \mathrm{PI}_m(\varepsilon) &:= I(\vec{X}_{t+m}, \ldots, \vec{X}_{t+1}; \vec{X}_t, \ldots, \vec{X}_{t-m+1}; \varepsilon) \\ &= 2H_m(\varepsilon) - H_{2m}(\varepsilon) \end{aligned} \tag{18}$$

If the limits (14) and (17) exist the predictive information $\mathrm{PI}(\varepsilon)$ is equal to the excess entropy $E(\varepsilon)$. For the finite time variants in general $\mathrm{PI}_m(\varepsilon) \neq E_m(\varepsilon)$:

$$\mathrm{PI}_m = E_m + \sum_{k=m}^{2m-1} (h_{m-1} - h_k) \geq E_m \ .$$

However, if the system is Markov of order p the conditional probabilities will only depend on the previous p time steps, $p(\vec{X}_t|X_{t-1}, \ldots, X_{t-\infty}) = p(\vec{X}_t|X_{t-1}, \ldots, X_{t-p})$, hence $h_m = h_{m-1}$ for $m > p$ and therefore $E = E_{p+1} = \mathrm{PI}_{p+1}$.

2.2 Decomposing the excess entropy for continuous states

In the literature [2, 5, 18] both the excess entropy and the predictive information were studied only for a given partition — usually a generating partition. Thus, using the excess entropy as a complexity measure for continuous valued time series has to deal with the fact that its value will be different for different partitions — even for different generating ones.

In (16) we have seen that the excess entropy for deterministic systems becomes infinite in the limit $\varepsilon \to 0$. The same applies to the predictive information (18). Moreover, we have seen that the increase of these complexity measures with decreasing ε is controlled by the attractor dimension of the system. Does this means that in the case of deterministic systems the excess entropy as a complexity measure reduces to the attractor dimension? Not completely. The constant in (16) reflects not only the scale of the signal, but also statistical dependencies or memory effects in the signal, in the sense that it will be larger if the conditional entropies converge slower towards the KS-entropy.

How can we separate the different contributions? We will start by rewriting Eq (15) as a sum. Using the conditional entropies (13) we get

$$E_m(\varepsilon) = \sum_{k=0}^{m-1} (h_k(\varepsilon) - h_{m-1}(\varepsilon)) \ .$$

Using the differences between the conditional entropies

$$\delta h_m(\varepsilon) := h_{m-1}(\varepsilon) - h_m(\varepsilon) \tag{19}$$

the excess entropy can be rewritten as

$$E_m(\varepsilon) = \sum_{k=1}^{m-1} k\delta h_k(\varepsilon) \ . \tag{20}$$

Note that the difference $\delta h_m(\varepsilon)$ is the conditional mutual information

$$\delta h_m(\varepsilon) = I(\vec{X}_t : \vec{X}_{t-m}|\vec{X}_{t-1}, \ldots, \vec{X}_{t-m+1}) \ .$$

It measures dependencies over m time steps that are not captured by dependencies over $m-1$ time steps. In other words, how much uncertainty about X_t can by reduced if in addition to the $m-1$ step past also the mth is taken into account. For a Markov process of order m the δh_k vanish for $k > m$. In this case the sum Eq (20) contains only a finite number of terms. On the other side truncating the sum at finite m could be interpreted as approximating the excess entropy by the excess entropy of an approximating Markov process. What can be said about the scale dependency of the $\delta h_m(\varepsilon)$? From Eqs. (2) and (3) follows that $h_m(\varepsilon) \propto -log(\varepsilon)$ for $m \leq D$ and $h_m(\varepsilon) \propto$ const for $m \geq D$. Using this and Eq. (19) we have to distinguish four cases for deterministic systems. Note that $\delta D = D - \lfloor D \rfloor$ denotes the fractal part of the attractor dimension.

$$\delta h_m(\varepsilon) \propto \begin{cases} \text{const} & \text{if} \quad m \leq D-1 \\ -(1-\delta D)\log\varepsilon & \text{if} \quad m = \lfloor D \rfloor \\ -\delta D \log\varepsilon & \text{if} \quad m = \lfloor D \rfloor + 1 \\ \text{const} & \text{if} \quad m > D+1 \ . \end{cases} \tag{21}$$

Thus the sum Eq. (20) can be decomposed into three parts:

$$E_m(\varepsilon) = \sum_{k=1}^{\lfloor D \rfloor - 1} k\delta h_k(\varepsilon) + \sum_{k=\lfloor D \rfloor}^{\lfloor D \rfloor + 1} k\delta h_k(\varepsilon) + \sum_{\lfloor D \rfloor + 2}^{m-1} k\delta h_k(\varepsilon) \ . \tag{22}$$

with the ε dependence showing up only in the middle term (MT):

$$\mathrm{MT} = \sum_{k=\lfloor D \rfloor}^{\lfloor D \rfloor + 1} k\delta h_k(\varepsilon) \propto c - \log \frac{\varepsilon}{\varepsilon_D} \tag{23}$$

with ε_D denoting the length scale where the deterministic scaling range starts. Therefore we have decomposed the excess entropy in three terms: Two ideally ε independent terms and one ε-dependent term. The first term

$$E_{\mathrm{state}} = \sum_{k=1}^{\lfloor D \rfloor - 1} k\delta h_k + c \tag{24}$$

will be called "state complexity" in the following because it is related to the information encoded in the state of the system. The constant c was added here in order to ensure that the ϵ-dependent term vanishes at ϵ_D — the beginning of the deterministic scaling range. The second ε-independent term

$$E_{\mathrm{mem}} = \sum_{\lfloor D \rfloor + 2}^{m-1} k\delta h_k(\varepsilon) \tag{25}$$

will be called "memory complexity" because it is related to the dependencies between the states on different time steps. What we call "state" in this context is related to the minimal embedding dimension to see the deterministic character of the dynamics which is $\lfloor D \rfloor + 1$ [11]. In order to be able to get a one to one reconstruction of the attractor a higher embedding dimension might be necessary [36]. Both ε-independent terms together we will call "core complexity"

$$E_{\mathrm{core}} = E_{\mathrm{state}} + E_{\mathrm{mem}} \,. \tag{26}$$

So far we only addressed the case of a deterministic scaling range. In the case of a noisy chaotic system we have to distinguish two ε regions: the deterministic scaling range described above and the noisy scaling range with $h_m(\varepsilon) \approx h_m^c - d \log \varepsilon$ with h_m^c determined by the noise level. In the stochastic scaling range all δh_m become ε-independent and the decomposition (22) seems to become unnecessary. This is not completely the case. Let us assume that the crossover between the two regions happens at a length scale ε^*. Moreover, let us assume that for sufficiently large m we have in the deterministic scaling range $h_m(\varepsilon) \approx h_{KS}$ (cf. Figs. 2(d)–(f)). Then we have

$$h_m(\varepsilon^*) \approx h_{KS} \approx h_m^c - \log \varepsilon^* \tag{27}$$

which allows to express the cross-over scale ε^* in terms of the KS-entropy and the noise level related continuous entropy h_m^c

$$\log \varepsilon^* \approx h_m^c - h_{KS} \,. \tag{28}$$

Moreover, the excess entropy in the *deterministic scaling range* will behave as

$$E_m(\varepsilon) \approx E_{state} + E_{mem} - D \log \frac{\varepsilon}{\varepsilon_D} \qquad \text{for} \quad \varepsilon \geq \varepsilon^* \,. \tag{29}$$

Evaluating this expressing at the crossover length scale ε^* allows to express the value of the excess entropy in the *stochastic scaling range* as

$$E_m \approx E_{state} + E_{mem} + D \log \varepsilon_D - D(h_m^c - h_{KS}) \qquad \text{for} \quad \varepsilon \leq \varepsilon^* \,. \tag{30}$$

In particular, this expression shows that in decreasing the noise level, i. e. h_m^c will increase the asymptotic value of the excess entropy for noisy systems. Thus, an increased excess entropy or predictive information for a fixed length scale or partition can be achieved in many ways:

1. by increasing dimension D of the dynamics
2. by decreasing the noise level ε^*
3. by increasing the amplitude ε_D
4. by increasing the state complexity
5. by increasing the correlations measured by the "memory" complexity, i.e. by increasing the predictability
6. by decreasing the entropy rate h_{KS}, i.e. by decreasing the unpredictability

Naturally, the effect of the noise level will be observed in the stochastic scaling range only. In practice there might be more than one deterministic and stochastic scaling range or even no clear scaling range at all. How we will deal with these cases will be described below when we introduce our algorithm.

2.3 Methods for estimating the information theoretic measures

Reliably estimating entropies and mutual information is very difficult in high-dimensional spaces due to the increasing bias of entropy estimates. Therefore we will employ two different approaches. On the one hand we will use an algorithm for the estimation of the mutual information proposed by **(author?)** [26] based on nearest neighbor statistics which allows to reduce the bias by employing partitions of different sizes in spaces of different dimensions. On the other hand we calculate a proxy for the excess entropy using correlation entropies [44] of order $q = 2$. These are related to the Rényi entropies of second order and the correlation sum provides an unbiased estimator. Both methods do not require binning but differ substantially in what they compute.

Estimation via local densities from nearest neighbor statistics (KSG) The most common approach to estimate information quantities of continuous processes, such as the mutual information, is to calculate the differential entropies (1) directly from the nearest neighbor statistics. The key idea is to use nearest neighbor distances [12, 25, 46] as proxies for the local probability density. This method corresponds in a way to an adaptive bin-size for each data point. For the mutual information $I(X, Y)$ (required e. g. to calculate the PI (18)), however, it is not recommended to naively calculate it directly from the individual entropies of X, Y and their joint (X, Y) because they may have very dissimilar scale such that the adaptive binning leads to spurious results. For that reason a new methods was proposed by **(author?)** [26], that we call *KSG*, which only uses the nearest neighbor statistics in the joint space. We denote $I^{(k)}_{\mathrm{KSG}}(X, Y)$ the mutual information estimate where k nearest neighbors where used for the local estimation.

The length scale on which the mutual information is estimated by this algorithm depends on the available data. In the limit of infinite amount of data $I^{(k)}_{\mathrm{KSG}}(X, Y) = \lim_{\varepsilon\to 0} I(X, Y, \varepsilon)$ for $k \gg 1$. However, in order to evaluate the quantity at a certain length scale (similar to ε above) and assuming the same underlying space for X and Y, noise of strength η is added to the data resulting in

$$I^{(k)}_{\mathrm{KSG}}(X, Y, \eta) = I^{(k)}_{\mathrm{KSG}}(X + U(\eta), Y + U(\eta)) \tag{31}$$

where $U(\eta)$ is the uniform distribution in the interval $[0, \eta]$. The idea of adding noise is to make the processes X and Y independent within neighborhood sizes below the length scale of the noise. In this way only the structures above the added noise-level contribute to the mutual information. Note that for small η the actual scale (k-neighborhood size) may be larger due to sparsity of the available data.

Estimation via correlation sum The correlation sum is one of the standard tools in nonlinear time series analysis [23, Chapter 6], [20]. Normally it is used to estimate the attractor dimension. However, it can also be used to provide approximate estimates of entropies and derived quantities such as the excess entropy. The correlation entropies for a random variable $\vec{X}$ with measure $\mu(\vec{x})$ are defined as [44]

$$H^{(q)}(\vec{X}; \varepsilon) = \begin{cases} -\frac{1}{q-1} \log \int \mu(B(\vec{x}, \varepsilon))^{q-1} d\mu(\vec{x}) & \text{for} \quad q \neq 1 \\ -\int \log \mu(B(\vec{x}, \varepsilon)) d\mu(\vec{x}) & \text{for} \quad q = 1 \end{cases} \tag{32}$$

where $B(\vec{x}, \varepsilon) = \{\vec{y} \mid ||\vec{y} - \vec{x}|| < \varepsilon\}$ is the "ball" at $\vec{x}$ with radius ε. For $q = 2$ (32) becomes $H^{(2)}(\vec{X}; \varepsilon) = -\log \int \mu(B(\vec{x}, \varepsilon)) d\mu(\vec{x})$. The integral in this formula is also known as "correlation integral". For N data points $\vec{x}_i \in \mathbb{R}^n \quad i = 1, \ldots, N$ it can be estimated using the correlation sum, which is the averaged relative number of pairs in an ε-neighborhood [23, Chapter 6],[20]

$$C^{(2)}(\varepsilon) = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \Theta\left(\varepsilon - \|\vec{x}_i - \vec{x}_j\|\right) . \tag{33}$$

Θ denotes the Heaviside function

$$\Theta(x) = \begin{cases} 0 & : \quad x < 1 \\ 1 & : \quad x \geq 1 . \end{cases}$$

Then the correlation entropy is

$$H^{(2)}(\varepsilon) = -\log C^{(2)}(\varepsilon) .$$

For sufficiently small ε it behaves as

$$H^{(2)}(\varepsilon) \approx \text{const} - D^{(2)} \log \varepsilon$$

with $D^{(2)}$ being the correlation dimension of the system [21]. A scale dependent correlation dimension can be defined as difference quotient

$$D^{(2)}(\varepsilon) = \log \frac{C^{(2)}(\varepsilon+\Delta)}{C^{(2)}(\varepsilon)} \left(\log \frac{\varepsilon+\Delta}{\varepsilon} \right)^{-1} . \tag{34}$$

For a temporal sequence of states (or state vectors) $x_i \quad i = 1, \ldots, m$ we can now define block entropies $H_m^{(2)}(\varepsilon) = -\log C^{(2)}(\varepsilon)$ by using $m \times d$-dimensional delay vectors $\vec{x}_i$. Now, we can define also the quantities corresponding to conditional entropies and to the excess entropy using the correlation entropy

$$h_m^{(2)}(\varepsilon) = H_{m+1}^{(2)}(\varepsilon) - H_m^{(2)}(\varepsilon) \tag{35}$$

$$E_m^{(2)}(\varepsilon) = H_m^{(2)}(\varepsilon) - m h_{m-1}^{(2)}(\varepsilon) . \tag{36}$$

We expect the same qualitative behavior regarding the ε-dependence of these quantities as for those based on Shannon entropies, see Eqs (11, 15). Quantitatively there might be differences, in particular for large ε and strongly non-uniform measures.

A comparison of the two methods with analytical results are given in the Appendix A, where we find a good agreement. Although, the KSG method seems to underestimate the mutual information for larger embeddings (higher dimensional state space). The correlation integral method uses a unit ball of diameter 2ε whereas the KSG method measures the size of the hypercube enclosing k neighbors where the data was subject to additive noise in the interval $[0, \eta]$. Thus comparable results are obtained with $\eta \approx 2\varepsilon$.

2.4 Algorithm for excess entropy decomposition

We are now describing the algorithm used to compute the proposed decomposition of the excess entropy in Sec. 2.2. The algorithm is composed of several steps: preprocessing, determining the middle term (MT) (23), determining the constant in MT, and the calculation of the decomposition and of quality measures.

Preprocessing: Ideally the δh curves are composed of straight lines in a log-linear representation, i. e. of the form $o - s \log(\varepsilon)$. We will refer to s as the slope (it is actually the inverted slope). Thus we perform fitting, that attempts to find segments following this form, details are provided in the Appendix B. Then the data is substituted by the fits in the intervals where the fits are valid. As for very small scales the δh become very noisy we extrapolate below the fit with the smallest scale. In addition we calculate the derivative $\hat{s}(m, \varepsilon) = \frac{d\, \delta h_m}{d\, \varepsilon}$ in each point, either from the fits (s, where available) or from finite differences of the data (using 5 points averaging).

Determining MT: In theory only two δh_m should have a non-zero slope at each scale ε, see Eq (21). However, in practice we have often more terms, such that we need to find for each ε the maximal range (m_l, m_u), where $\forall i \in [m_l, m_u] : \hat{s}(i, \varepsilon) > s_{\min}$, i. e. the slope is larger than the threshold $s_{\min}$. However, this is only valid for deterministic scaling ranges. In stochastic ranges all δh should have zero slope. We introduce a measure of stochasticity, defined as $\kappa(m, \varepsilon) = 1 - \sum_{k=m_l}^{m_u} \hat{s}(k, \varepsilon)$ which is 0 for purely deterministic ranges and 1 for stochastic ones. The separation between state and memory complexity is then inherited from the next larger deterministic range. Thus if $\kappa(m, \varepsilon) \geq \kappa_{\max}$ we use (m_l, m_u) at ε^*, where $\varepsilon^* = \arg\min_{e \in (\varepsilon, \infty)} \kappa(m, e) < \kappa_{\max}$. Note that the here algorithmically defined ϵ^* is not necessarily equal to the ϵ^* defined above (28) for an ideal-typical noisy deterministic system.

Determining the constant in MT: In order to obtain the scale-invariant constant c of the MT, see Eq (23), we would have to define a certain length scale ε_D. Since this cannot be done robustly in practice (in particular because it may not be the same ε_D for each m) we resort to a different approach. The constant parts of the $\delta h_{m_l \ldots m_u}$ terms in the MC can be determined from plateaus on larger scales. Thus, we define $c_m^{\mathrm{MT}}(\varepsilon) = \min_{e \in (\varepsilon, \varepsilon^*]} \delta h_m(e)$, where ε^* is smallest scale $> \varepsilon$ where we have a near-zero slope, i. e.. $\varepsilon^* = \arg\min_{e \in (\varepsilon, \infty)} s(m, e) < s_{\min}$. In case there is no such ε^* then $c_m^{\mathrm{MT}} = 0$.

Decomposition and quality measures: The decomposition of the excess entropy follows Eqs (24, 25) with m_l and m_u used for splitting the terms:

$$E_{\mathrm{state}} = \sum_{k=1}^{m_l - 1} k \delta h_k + \sum_{k=m_l}^{m_u} k c_k^{\mathrm{MT}} \qquad E_\varepsilon = \sum_{k=m_l}^{m_u} k(\delta h_k - c_k^{\mathrm{MT}}) \qquad E_{\mathrm{mem}} = \sum_{k=m_u+1}^{m} k \delta h_k \tag{37}$$

quantity	definition	description
κ=stochastic(ε)	$1 - \sum_{k=m_l}^{m_u} \hat{s}(k,\varepsilon)$	0: fully deterministic, 1: fully stochastic at ε
% negative(ε)	$1/m \sum_{k=1}^{m} [\![\delta h(k,\varepsilon) < 0]\!]$	percentage of negative δh
% no fits(ε)	$1/m \sum_{k=1}^{m} [\![\text{valid fit for } \delta h(k,\varepsilon)]\!]$	percentage of δh where no fit is available
% extrap.(ε)	$1/m \sum_{k=1}^{m} [\![\text{is } \delta h(k,\varepsilon) \text{ extrapolated}]\!]$	percentage of δh that where extrapolated

Table 1: Quality measures for decomposition algorithm. We use the Iverson bracket for Boolean expression: $[\![\text{True}]\!] := 1$ and $[\![\text{False}]\!] := 0$. They are all normalized to $[0, 1]$ where typically 0 is the best score and 1 is the worst.

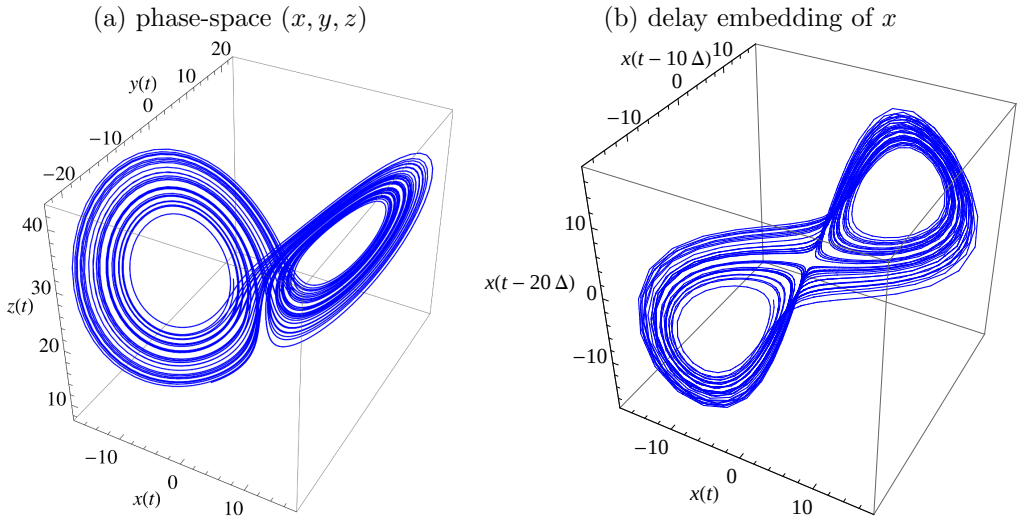

Figure 1: Lorenz attractor with standard parameters. **(a)** Trajectory in state space (x, y, z). **(b)** Embedding of x with $m = 3$, time-delay $\tau = 10$, and $\Delta = 0.01$.

In addition we can compute different quality measures to indicate the reliability of the results, see Table 1.

2.5 Illustrative example

To get an understanding of the quantities, let us first apply the methods to the Lorenz attractor, as a deterministic chaotic system, and to its noisy version as a stochastic system.

Deterministic system: Lorenz attractor The Lorenz attractor is obtained as the solution to the following differential equations:

$$\dot{x} = s(-x + y) \tag{38}$$
$$\dot{y} = -xz + rx - y \tag{39}$$
$$\dot{z} = -xy - bz \tag{40}$$

integrated with standard parameters $s = 10$, $r = 28$, $b = 8/3$. Fig 1 displays the trajectory in phase-space (x, y, z) and using delay embedding of $x : (x(t), x(t - \tau\Delta), \ldots, x(t - m\tau\Delta))$ using $m = 3$ and time delay is $\tau = 10$ (for sampling time $\Delta = 0.01$). The equations where integrated with Runge-Kutta fourth order with step size $5 \cdot 10^{-4}$. An alternative time delay may be obtained from first minimum of mutual information as a standard choice [15, 30], which would be $\tau = 18$, but $\tau = 10$ yields clearer results for presentation.

The result of the correlation sum method using TISEAN [22] is depicted in Fig 2 for 10^6 data points. From the literature we know that the chaotic attractor has a fractal dimension of 2.06 [21] which we can verify with the correlation dimension $D_m^{(2)}$, see Fig 2(a) for $m > 2$ and small ε. The conditional $h_m^{(2)}$ (d) becomes constant and identical for larger m. The excess entropy $E_m^{(2)}$ (g) approaches the predicted scaling behavior of: const $- D \log(\varepsilon)$ (16) with $D = 2.06$ and const $= 4.75$. The predictive information PI_m (j) shows the same scaling behavior on the coarse range but with a smaller slope and saturates at small ε.

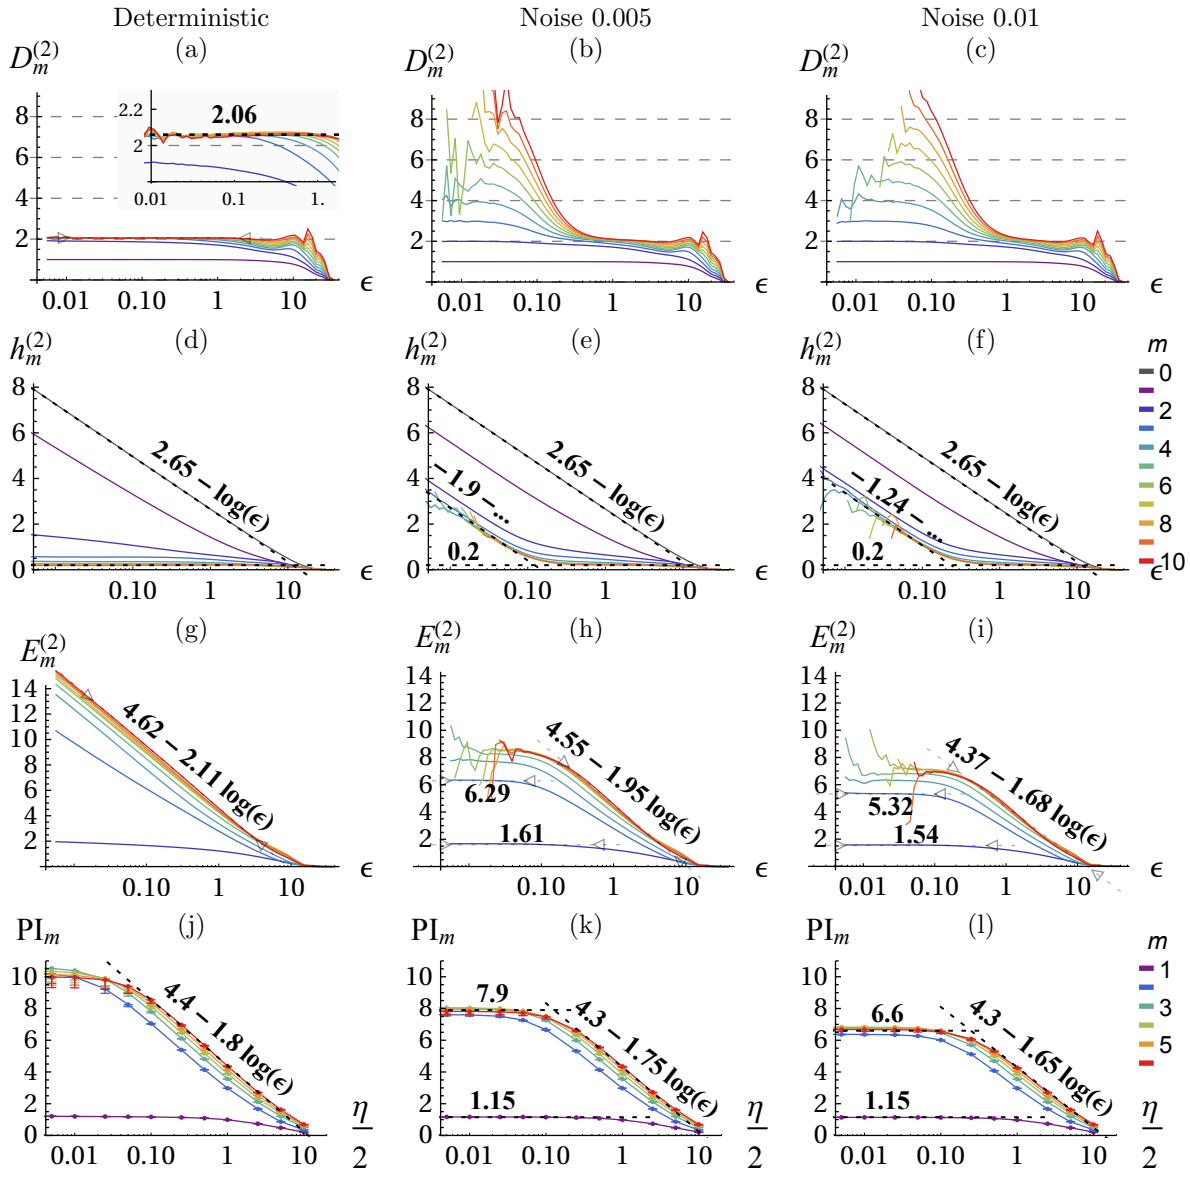

Figure 2: Correlation dimension $D_m^{(2)}$ **(a)–(c)** (34), conditional block entropies $h_m^{(2)}$ **(d)–(f)** (35), excess entropy $E_m^{(2)}$ **(g)–(i)** (36), and predictive information PI_m **(j)–(l)** (18) of the Lorenz attractor estimated with the correlation sum method (a)–(i) and KSG (j)–(l) for the deterministic system (first column) and with absolute dynamic noise 0.005, and 0.01 (second and third column respectively). The error bars in (j)–(l) show the values calculated on half of the data. All quantities are given in nats (natural unit of information with base e) and in dependence of the scale (ε,η in space of x Eq (38)) for a range of embeddings m, see color code. In (d)–(f) the fits for $h_0 = H_1$ allow to determine $H^C = 2.65$ and for h_m give h_{KS} and h_m^c, Eqs (27–30). Parameters: delay embedding of x with $\tau = 10$.

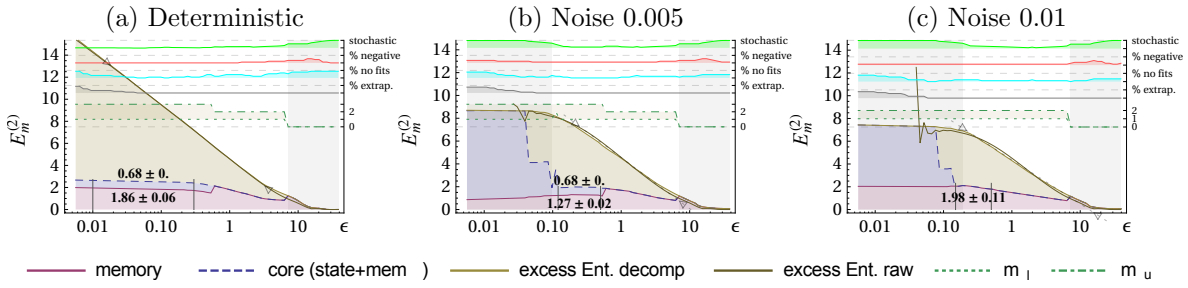

Figure 3: Excess Entropy decomposition for the Lorenz attractor into state complexity (blue shading), memory complexity (red shading), and ε-dependent part E_ε (beige shading) (all in nats). Columns as in Fig 2. A set of quality measures and additional information is displayed on the top using the right axis. m_l and m_u refer to the terms in E_ε (37). Scaling ranges that are identified as stochastic are shaded in gray (stochastic indicator $> \kappa_{\max} = 0.5$). In manually chosen ranges, marked with vertical black lines, we evaluate the mean and standard deviation of the memory- and state complexity. Parameters: $s_{\min} = 0.1$.

		determ.	noise 0.005	noise 0.01	noise 0.02
Eq (16)	D	2.11	1.95	1.68	1.59
Eq (16)	const	4.62	4.55	4.37	4.21
Eq (27)	ε^*	0	0.12	0.23	0.45
Eq (24)	E_{state}	0.68±0	0.68±0	0	0
Eq (25)	E_{mem}	1.86±0.06	1.27±0.02	1.98±0.11	1.2±0.1
Eq (26)	E_{core}	2.54±0.06	1.95±0.02	1.98±0.11	1.2±0.1

Table 2: Decomposition of the excess entropy for the Lorenz system. Presented are the unrefined constant and dimension of the excess entropy, the state-, memory-, and core complexity (in nats) determined in the deterministic scaling range, see Fig 3.

Stochastic system: noisy Lorenz attractor In order to illustrate the fundamental differences between deterministic and stochastic systems when analyzed with information theoretic quantities we consider now the Lorenz dynamical system with dynamic noise (additive noise to the state (x, y, z) in Eqs (38–40) before each integration step) as provided by the TISEAN package[22]. The dimension of a stochastic systems is infinite, i. e. for embedding m the correlation integral yields the full embedding dimension as shown in Fig 2(b),(c) for small ε.

On large scales the dimensionality of the underlying deterministic system can be revealed to some precision either from $D^{(2)}$ or from the slope of $E^{(2)}$ or PI. Thus, we can identify a determinstic and a stochastic scaling range with different qualitative scaling behavior of the quantities. In contrast to the deterministic case, the excess entropy converges in the stochastic scaling range and no more contribution from $m > 5$ are observed as shown in Fig 2(h),(i). By measuring the ε-dependence (slope) we can actually differentiate between deterministic and stochastic systems and scaling ranges, which is performed by the algorithm (Sec. 2.4) and presented in Fig 3. The predictive information (k)–(l) again yields a lower slope, meaning a lower dimension estimate, but is otherwise consistent. However, it does not allow to safely distinguish determinstic and stochastic systems, because it always saturates due to the effect of finite amount of data.

Let us now consider the decomposition of the excess entropy as discussed in Sec. 2.2 and Sec. 2.4 to see whether the method is consistent. Fig 3 shows the scale-dependent decomposition together with the determined stochastic and deterministic scaling ranges. The resulting values for comparing the determinstic scaling range are given in Table 2 and reveal, that the E_{core} decreases for increasing noise level similarly as the constant reduces. This is a consistency check because the intrinsic scale (amplitudes) of the deterministic and noisy Lorenz system are almost identical. We also checked the decomposition for a scaled version of the data where we found the same values state-, memory- and core complexity. We see also that the state complexity E_{state} vanishes for larger noise, because the dimension drops significantly below 2, where we have no summands in the state complexity and the constant c^{MT} is also 0.

To summarize, in deterministic systems with a stochastic component, we can identify different ε-ranges (scaling ranges) within which stereotypical behaviors of the entropy quantities are found. This allows us to determine the dimensionality of the deterministic attractor and decompose the excess entropy in order to obtain two ε-independent complexity measures.

Robot	Behavior	Length	Dataset	Video
SNAKE	side rolling	$9 \cdot 10^5$	D-S1	Video S1
SNAKE	crawling	$9 \cdot 10^5$	D-S2	Video S2
HEXAPOD	jumping after 2 min	10^6	D-H1	Video H1
HEXAPOD	jumping after 4 min	10^6	D-H2	Video H2
HEXAPOD	jumping after 8 min	10^6	D-H3	Video H3

Table 3: Experiments, data sets and videos. See playfulmachines.com/QuantBeh2015.

3 Results

In this section we will apply our methods to the behavior of robotic systems. We use a previously published control scheme that can generate a large variety of smooth and coherent behaviors from a local learning rule [6, 8]. As it is a simplified version of the time-local predictive information maximization controller we aim at answering the hypothesis that the controller actually increases the complexity of behavior during learning.

3.1 Application to robotics: controller and learning

Let us now briefly introduce the control framework and the robots used to generate the behavior we are analyzing in the following. The robots are simulated in a physically realistic rigid-body simulation [29]. Each robot has a number of sensors producing real-valued stream $s \in \mathbb{R}^n$ and a number of actuators, also controlled by real values $a \in \mathbb{R}^o$. The sensors measure the actual joint angles and the body accelerations. In all examples the actuators are position controlled servo motors with the particular property to have no power around the set-point. This enables small perturbations to be perceivable in the joint-position sensor values. The time is discretized with an update frequency of 50 Hz.

The control is done in a reactive closed-loop fashion by a one layer feed-forward neural network:

$$a_t = K(s_t, C, h) = \tanh(Cs_t + h), \tag{41}$$

where C is a weight matrix, h is a weight vector, and tanh is applied component-wise. For fixed weights this controller is very simple, however it can generate non-trivial behavioral modes due to the complexity of the body-environment interaction. For instance with a scaled rotation matrix C an oscillatory behavior may be controlled which frequency is adaptive to how fast the physical system follows. To get a variety of behavioral modes the parameters (C, h) have to be changed which can be done by optimizing certain objective functions. Several methods have been proposed that use generic (robot independent) internal drives. Homeokinesis [7, 10] for instance balances dynamical instability with predictability. More recently the maximization of the time-local predictive information of the sensor process was used as driving force [28]. On an intuitive level the maximization of the PI (18) of the sensor process leads to a high variety of sensory input due to the entropy term $H(S)$ while keeping a temporal structure by minimizing the future uncertainty reflected in the conditional entropy $H(S'|S)$. Indeed we have found that it produces smooth and coordinated behavior in different robotic systems, including hexapod robots and humanoids. On the other hand, maximizing PI also leads to behavior of high complexity with maximal dimension as discussed in Sec. 2.2. A simplified version of the PI-maximizing algorithm, called ULR published in [6, 8], leads to more coherent and defined behavioral modes of seemingly low dimensionality which we use in this paper. In order to apply the method to behavior of robots or possibly to that of animals the state space reconstruction has to be done from certain measurements. Here we use a single sensor value, but the formalism can be extended to multiple sensor values as well. It must be guaranteed that sufficient physical coupling exist between the measured quantity and the rest of the body.

3.2 Experiments

Let us now present the emergent behaviors obtained from the control algorithm ULR. We use two robots, a snake-like robot and an insect-like hexapod robot as illustrated in Fig 4. The behaviors of the snake are published here for the first time, whereas the behaviors of the hexapod have been published before in **(author?)** [8]. However, this paper focuses primarily on the quantification of behavior and not on its generation. The SNAKE has identical cylindrical segments which are connected by actuated 2 degrees of freedom (DoF) universal joints. In this way two neighboring segments can be inclined to each other, but without relative torsion along the segment-axes. The sensor values are the actual positions of the joints. If the robot is controlled with the ULR controller very quickly highly coordinated behavioral

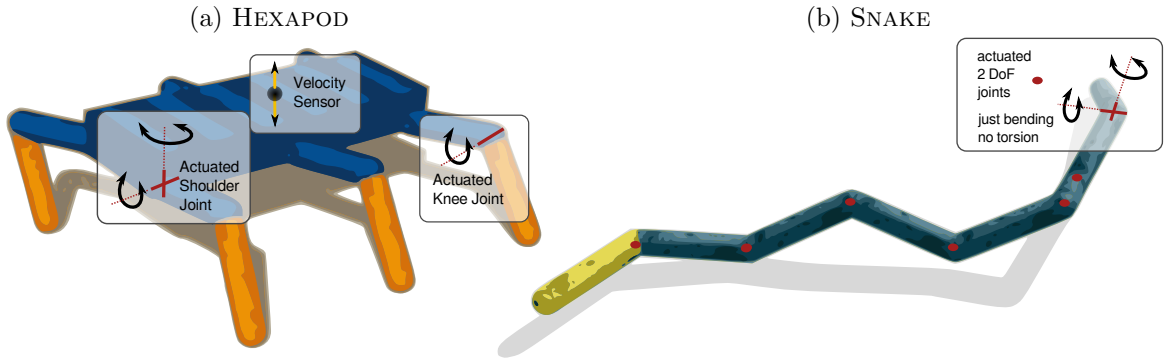

Figure 4: Robots used for the experiments. **(a)** The Hexapod. 18 actuated DoF: 2 shoulder and 1 knee joint per leg. **(b)** The Snake. 14 DoF: 8 segments connected by 2 DoF joints each.

Figure 5: Side-rolling and crawling of the Snake. Note that also for rolling each actuator has to act accordingly because no torsion is possible, see Fig 4(b).

modes emerge, such as crawling or side-rolling, see Fig 5. An external perturbation can bring the robot into another behavior. Without perturbation the robot typically transition into one behavior and remain in it. We recorded $9 \cdot 10^5$ samples from a typically side-rolling and crawling behavior.

The Hexapod is a insect-like robot with 18 DoF, see Fig 4(a). The sensor values are the joint angles and the vertical velocity of the trunk. Note, that the actual phase space dimension is $18 + 18 + 4 = 40$ as we have joint angles, joint velocities, and height and orientation of the trunk. Controlled with the ULR controller the robot develops a jumping behavior, see Fig 6, within 8 min of simulated time [8]. The behavior of the robot changes due to the update dynamics of the controller. For the analysis a stationary process is required, such that we stop the update at some time and obtain a static controller (fixed C, h). Using this the robots behavior stays essentially the same. We analyze three behaviors, where the learning was stopped at 2 min, 4 min and 8 min after the start and recorded each 10^6 samples. After "2 min" the robot moves periodically up and down. The "4 min" behavior shows already signs of jumping and the final behavior after "8 min" is a whole body periodic jumping as displayed in Fig 6.

3.3 Quantifying the Behavior

We start with analyzing the data from the Snake experiments followed by the Hexapod experiments.

3.3.1 Snake

In Fig 7 the result of the quantification of the two behaviors (Fig 5) are presented. We used the delay embedding of the sensor x_7 which is one of the joint angular sensors of the middle joint. We find that the side-rolling behavior is roughly 1-dimensional on the behaviorally relevant scales ($0.05 < \varepsilon < 0.5$), see Fig 7(c). On smaller scales the dimension raises with each embedding although the embedding dimension is not reached on the observable length scales and therefore the excess entropy does not saturate. On the larger scales we find the expected scaling behavior (16) with $D \approx 1$.

The crawling behavior is also 1-dimensional on the coarse length scales ($0.1 < \varepsilon < 0.5$), see Fig 7(d). The smaller scales, however, bear a surprise because the dimension raises to a value of ≈ 2.5. The excess

Figure 6: Jumping motion pattern emerging with vertical speed sensor 8 min after the start.

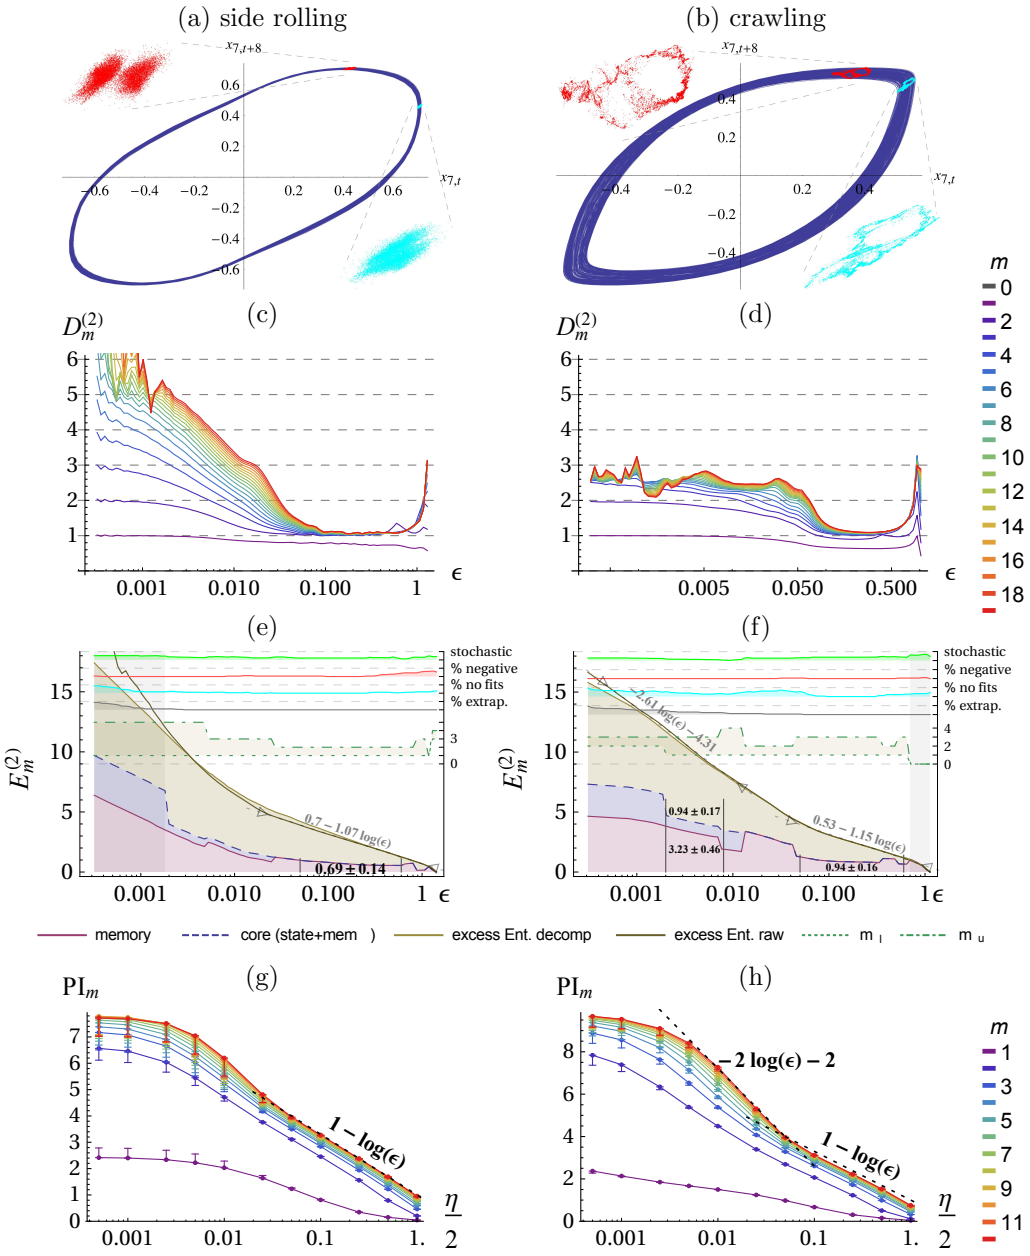

Figure 7: Quantification of two SNAKE behaviors. First column: side rolling; second column: crawling (see Fig 5). **(a),(b)** Phase plots with Poincaré sections at the maximum of first and second embedding. **(e),(f)** Excess entropy with decomposition (in nats), see Fig 3 for more details. **(g),(h)** Predictive information (in nats), see Fig 2 for more details. Parameters: delay embedding of sensor x_7 (middle segment) with $\tau = 8$, $s_{\min} = 0.05$, $\kappa = 0.5$.

Table 4: Decomposition of the excess entropy for the SNAKE. Dimension and unrefined constant, state-, memory-, and core complexity (in nats) on the coarse deterministic scaling range $\varepsilon \in (0.05, 0.6)$.

		side rolling	crawling
Eq (16)	D	1.07	1.15
Eq (16)	const	0.7	0.53
Eq (24)	E_{state}	0.0±0	0±0
Eq (25)	E_{mem}	0.69±0.14	0.94±0.16
Eq (26)	E_{core}	0.69±0.14	0.94±0.16

entropy shows also a slope of $D \approx 1.0$ on the coarse length scales and a slope of $D \approx 2.5$ on the smaller scales, see Fig 7(f).

Calculating the predictive information using the KSG method yields a similar result on the large scales, as shown in Fig 7(g),(h). For smaller scales, however, the predictive information always saturates, as seen before in Fig 2(j). In the crawling case we get again an underestimated dimension of $D \approx 2$.

The decomposition of the excess entropy follows for the above experiments are given in Table 4. Interestingly, the E_{core} attributes a significantly higher complexity to the crawling behavior, whereas the plain excess entropy constant would suggest that side-rolling is more complex. Since we take out the effect of different scales our new quantities are more trustworthy.

Note, that for the crawling behavior at the small scaling range ($\varepsilon < 0.01$) we can also perform the decomposition and obtain $E_{\text{state}} = 0.84 \pm 0.17$, $E_{\text{mem}} = 3.23 \pm 0.46$, however these values are mostly useful for comparison, which we only have on the large scale. Note that $E_{\text{state}}$and $E_{\text{mem}}$are not really constant in the plots because the deterministic scaling range is not very good in this case (cf. the quality measure "no fits" in Fig. 7(f)).

3.3.2 Hexapod

Fig 8 shows the quantification of the three successive behaviors of the HEXAPOD robot, reconstructed from the delay embedding of sensor x_1 (up-down direction of right hind leg). We clearly see an increase in dimensionality of the behavior, Fig 8(d)–(f) which we can expect because the high jumping behavior requires more compensation movements due to different landing poses. For the first behavior we can identify three scaling ranges of different dimensionality and complexity. On the coarse level $[0.1, 0.6]$ the behavior is described by a dimension of 1.2 whereas on medium scales $[0.006, 0.02]$ we find an approximately 1.6 dimensional behavior and on the small scales $[0.001, 0.005]$ it appears 5 dimensional.

For the second behavior we find only a short plateau in Fig 8(e) at $[0.06, 0.3]$ of dimensionality 4. The excess entropy shows a lower dimensionality of ≈ 3.4, see Fig 8(h). For the third behavior we cannot identify a clear plateau in the dimension plot. The excess entropy allows us to attribute a dimensionality and core complexity to the behavior on large scales, as summarized in Table 5. Whereas the scaling behavior of the excess entropy suggests a similar dimensionality, the core complexity suggest that the 8 min behavior is more complex.

The predictive information calculated with the KSG method underestimates the dimension by far and is only able to report structure on the coarsest scale, see Fig 8(j)–(l) and compare for instance slope 1.8 of PI in (k) with slope 3.36 of $E^{(2)}$ in (h). Note that usually only the value of the mutual information without added noise (or only very small one) is reported (left-most value in our graphs). In our case this would attribute the highest complexity to the 2 min behavior and the lowest to the 8 min behavior, see also Table 5. This is because the intrinsic noise level increases from 2 min to 8 min, such that the plateau level is much lower.

4 Discussion

In recent years research in autonomous robots has been more and more successful in developing algorithms for generating behavior from generic task-independent objective functions [1, 14, 24, 28, 32, 35, 39, 48]. This is valuable, on the one hand, as a step towards understanding and creating autonomy in artificial beings, and on the other hand, to obtain eventually more effective learning algorithms for mastering real-life problems. This is because the task-agnostic drives aim at efficient exploration of the agent's possibilities by generically useful behavior [7, 28, 31, 32, 38] and at generating an operational state of the system [13] from which specific tasks can quickly be accomplished.

Very often emerging behaviors are not the result of optimizing a global function but for instance self-organize from local interaction rules. We propose to use information theoretic quantities for quantifying

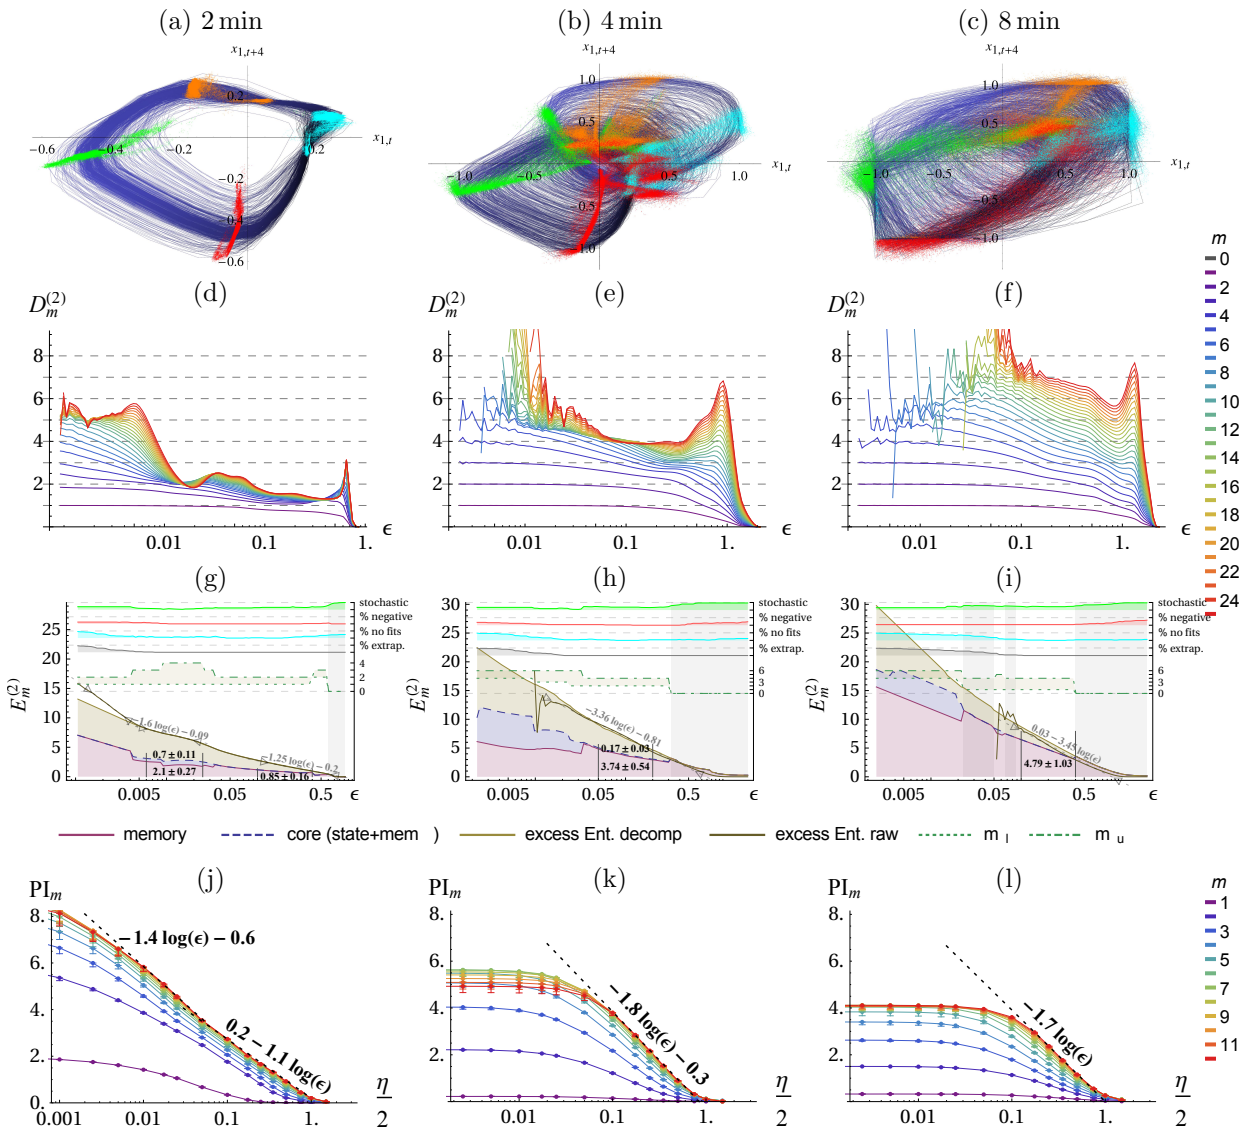

Figure 8: Quantification of the three HEXAPOD behaviors. Behaviors: first column: 2 min; second column: 4 min; third column 8 min, see Fig 6. **(a)–(c)** Phase plots with Poincaré sections at the maximum and minimum of first and second embedding. The color of the trajectory (blue–black) encodes for the third embedding dimension. **(g)–(i)** Excess entropy with decomposition (in nats), see Fig 3 for more details. The excess entropy fit at the smallest scale in (g) is $-18 - 5\log(\varepsilon)$. **(j)–(l)** Predictive information (in nats), see Fig 2 for more details. Parameters: delay embedding of sensor x_1 (up-down direction of right hind leg) with $\tau = 4$, $s_{\min} = 0.05$, $\kappa_{\max} = 0.5$.

		2 min (fine)	2 min (coarse)	4 min (coarse)	8 min(coarse)
	ε-range	$(0.006, 0.025)$	$(0.1, 0.4)$	$(0.05, 0.2)$	$(0.1, 0.4)$
Eq (16)	D	1.6	1.25	3.36	3.45
Eq (16)	const	-0.09	-0.2	-0.81	0.03
Eq (24)	E_{state}	0.7±0.11	0±0	0.17±0.03	0±0
Eq (25)	E_{mem}	2.1±0.27	0.85±0.16	3.74±0.54	4.79±1.03
Eq (26)	E_{core}	2.8±0.38	0.85±0.16	3.91±0.57	4.79±1.03
Eq (18)	$\text{PI}(\eta = 0)$	≈ 9		≈ 5.5	≈ 4

Table 5: Decomposition of the excess entropy for the HEXAPOD behaviors on the fine and coarse scale. For the 4 min and 8 min behavior we have no reliable estimate on the fine scale. The core complexity and the dimension rises from 2 min to 4 min. From 4 min to 8 min the dimension stayed the same but the core complexity increases. The last line provides the predictive information on the smallest scale.

these behaviors. There are already studies using information theoretic measures to characterize emergent behavior of autonomous robots [27, 40, 47] — in particular with the aim to use them as task-independent objective functions [1, 24, 28, 35, 39, 48]. Usually these information theoretic quantities are estimated for a fixed partition, i. e. a specific discretization, or by using differential entropies assuming that the data were sampled from a probability density. In this paper we show that it is both necessary and fruitful to explicitly take into account the scale dependence of the information measures.

We quantify and compare behaviors emerging from self-organizing control of high-dimensional robots by estimating the excess entropy and the predictive information of their sensor time series. The estimation is done on the one hand using entropy estimates based on the correlation integral and on the other hand the predictive information is directly estimated as a mutual information using an estimator (KSG) based on differential entropies. The latter implicitly assumes stochastic data, i. e. the existence of a smooth probability density for the temporal joint distribution, such that the predictive information converges towards a finite value. However, this is not the case for deterministic systems. Due to the adaptivity of the algorithm the predictive information is then calculated at a certain length scale which depends on the available amount of data. Thus, a naive application of the KSG estimator can lead to totally meaningless results, because it does not allow to control for the length scale. To avoid this problem and to control the length scale, we applied additive noise. Nevertheless, our results show that the complexity estimates based on the KSG estimator are much less reliable than those based on the correlation integral. While for simple stochastic systems such as autoregressive models both estimators deliver consistent results for small embedding dimensions, the predictive information is underestimated due to finite sample effects for larger embedding dimensions (cf. Fig 9). For more complex systems such as the Lorenz attractor or the behavioral data from robots the KSG algorithm underestimates the dimensionality, see e. g. Fig 8(j)–(l), does not allow to extract length scales properly, and blurs the distinction between stochastic and deterministic scaling ranges.

An important result of the presented study is that the complexity of the behavior can be different at different length scales: an example is the crawling behavior of the Snake which is 1-dimensional on the coarse scale and 2.6-dimensional on the small scales, see Fig 7 (the same for the Hexapod). In order to quantify these effects we analyzed for the first time the resolution dependence of the well-known complexity measure for time series known as excess entropy [5, 41], effective measure complexity [18], or predictive information [2]. We show that the scale dependent excess entropy reflects different properties of the system: scale, dimension and correlations. For deterministic systems the attractor dimension controls how the complexity increases with increasing resolution, while the offset depends on the overall scale of the behavior, e. g. the amplitude of an oscillation, and on the structure of the temporal correlations. In order to disentangle the latter two effects we introduced a new decomposition of the excess entropy with the *core complexity* quantifying the structural complexity beyond the dimension. In order to estimate it one has to remove the effect of the overall scale of the behavior from the excess entropy. We devised an algorithm based on an automatic identification of scaling ranges in the conditional mutual information $\delta h_m(\epsilon)$.

For noisy systems the excess entropy remains finite. The actual value, however, depends on the noise level, which is consistent with the intuition: The larger the noise the lower the complexity if all other factors remain constant. While we have already a good understanding of these quantities for systems with clearly identifiable scaling regions as in the noisy Lorenz system, very often real systems do not show such ideal types of behavior. Therefore we had to use heuristic procedures which provide plausible results for the data under study. Our quantification of the behaviors of the Snake found that crawling has a higher complexity than side-winding, even though the original excess entropy would have suggested the opposite. For the Hexapod as a learning system (aiming at maximizing time local PI) we were able to quantify the learning progress: The decomposition of the excess entropy revealed the strategy used to increase the complexity: (1) The dimensionality of the behavior increased on the large scales. This does not necessarily mean that the attractor dimension in the mathematical sense increased, because the latter is related to behavior on infinitely small scales which might be irrelevant for the character of the behavior. (2) The core complexity is increased on the large scales indicating also more complex temporal correlations in the learned behavior. We also discuss further strategies for increasing the complexity of a time-series, that could be employed in designing future learning algorithms, most notably, the decrease in the entropy rate and the increase of the state complexity.

Our study contributes a well grounded and practically applicable quantification measure for the behavior of autonomous systems as well as time-series in general. In addition, a set of quality indicators are provided for self-assessment in order to avoid spurious results.

Future work is to develop a better theoretical understanding of the scale dependence of the excess entropy for realistic behaviors that are not low dimensional deterministic or linearly stochastic such as

the behavior observed in the HEXAPOD after 8 minutes of learning, Fig 8. We expect also further insights by applying the same kind of analysis to movement data from animals and humans.

Acknowledgments

This work was supported by the DFG priority program 1527 (Autonomous Learning) and by the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 318723 (MatheMACS) and from the People Programme (Marie Curie Actions) of the European Union's Seventh Framework Programme (FP7/2007-2013) under REA grant agreement no. 291734.

References

[1] N. Ay, N. Bertschinger, R. Der, F. Güttler, and E. Olbrich. Predictive information and explorative behavior of autonomous robots. *The European Physical Journal B*, 63(3):329–339, 2008.

[2] W. Bialek, I. Nemenman, and N. Tishby. Predictability, complexity and learning. *Neural Computation*, 13:2409, 2001.

[3] A. Cohen and I. Procaccia. Computing the Kolmogorov entropy from time signals of dissipative and conservative dynamical systems. *Phys. Rev. A*, 31:1872–1882, 1985.

[4] T. M. Cover and J. A. Thomas. *Elements of information theory*. John Wiley & Sons, 2006.

[5] J. P. Crutchfield and D. P. Feldman. Regularities unseen, randomness observed: Levels of entropy convergence. *Chaos*, 13(1):25–54, 2003.

[6] R. Der. On the role of embodiment for self-organizing robots: behavior as broken symmetry. In M. Prokopenko, editor, *Guided Self-Organization: Inception*, volume 9 of *Emergence, Complexity and Computation*, pages 193–221. Springer, 2014.

[7] R. Der and G. Martius. *The Playful Machine - Theoretical Foundation and Practical Realization of Self-Organizing Robots*. Springer, Berlin Heidelberg, 2012.

[8] R. Der and G. Martius. Behavior as broken symmetry in embodied self-organizing robots. In *Advances in Artificial Life, ECAL 2013*, pages 601–608. MIT Press, 2013.

[9] R. Der and G. Martius. Novel plasticity rule can explain the development of sensorimotor intelligence. *to appear in PNAS*, [arXiv:abs/1505.00835], 2015.

[10] R. Der, G. Martius, and F. Hesse. Let it roll – emerging sensorimotor coordination in a spherical robot. In L. M. Rocha, L. S. Yaeger, M. A. Bedau, D. Floreano, R. L. Goldstone, and A. Vespignani, editors, *Proc, Artificial Life X*, pages 192–198. Intl. Society for Artificial Life, MIT Press, August 2006.

[11] M. Ding, C. Grebogi, E. Ott, T. Sauer, and J. A. Yorke. Plateau onset of correlation dimension: When does it occur? *Phys. Rev. Lett.*, 70:3872–3875, 1993.

[12] R. Dobrushin. A simplified method of experimentally evaluating the entropy of a stationary sequence. *Theory of Probability & Its Applications*, 3(4):428–430, 1958.

[13] S. Doncieux and J.-B. Mouret. Beyond black-box optimization: a review of selective pressures for evolutionary robotics. *Evolutionary Intelligence*, 7(2):71–93, 2014.

[14] M. Frank, J. Leitner, M. Stollenga, A. Förster, and J. Schmidhuber. Curiosity driven reinforcement learning for motion planning on humanoids. *Frontiers in Neurorobotics*, 7(25), 2014.

[15] A. M. Fraser and H. L. Swinney. Independent coordinates for strange attractors from mutual information. *Phys. Rev. A*, 33(2):1134–1140, 1986.

[16] P. Gaspard and X.-J. Wang. Noise, chaos and (ϵ, τ)-entropy per unit time. *Physics Reports*, 235:291–343, 1993.

[17] A. Giffin. From physics to economics: An econometric example using maximum relative entropy. *Physica A: Statistical Mechanics and its Applications*, 388(8):1610–1620, 2009.

[18] P. Grassberger. Toward a quantitative theory of self-generated complexity. *International Journal of Theoretical Physics*, 25(9):907–938, 1986.

[19] P. Grassberger. Randomness, information, and complexity. In F. Ramos-Gómez, editor, *Proceedings of the 5th Mexican School on Statistical Physics (EMFE 5)*, Singapore, 1991. World Scientific. corrected version arXiv:1208.3459.

[20] P. Grassberger. Grassberger-procaccia algorithm. *Scholarpedia*, 2(5):3043, 2007.

[21] P. Grassberger and I. Procaccia. Characterization of strange attractors. *Physical Review Letters*, 50(5):346, 1983.

[22] R. Hegger, H. Kantz, and T. Schreiber. Practical implementation of nonlinear time series methods: The tisean package. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 9(2):413–435, 1999.

[23] H. Kantz and T. Schreiber. *Nonlinear time series analysis*, volume 7. Cambridge University Press, 2004.

[24] A. S. Klyubin, D. Polani, and C. L. Nehaniv. Empowerment: a universal agent-centric measure of control. In *Congress on Evolutionary Computation*, pages 128–135, 2005.

[25] L. Kozachenko and N. N. Leonenko. Sample estimate of the entropy of a random vector. *Problemy Peredachi Informatsii*, 23(2):9–16, 1987.

[26] A. Kraskov, H. Stögbauer, and P. Grassberger. Estimating mutual information. *Phys. Rev. E*, 69:066138, Jun 2004.

[27] M. Lungarella and O. Sporns. Mapping information flow in sensorimotor networks. *PLoS Comput Biol*, 2(10):e144, 10 2006.

[28] G. Martius, R. Der, and N. Ay. Information driven self-organization of complex robotic behaviors. *PLoS ONE*, 8(5):e63400, 2013.

[29] G. Martius, F. Hesse, F. Güttler, and R. Der. LpzRobots: A free and powerful robot simulator. http://robot.informatik.uni-leipzig.de/software, 2010.

[30] E. Olbrich and H. Kantz. Inferring chaotic dynamics from time-series: on which length scale determinism becomes visible. *Physics Letters A*, 1997.

[31] P.-Y. Oudeyer, F. Kaplan, and V. Hafner. Intrinsic motivation systems for autonomous mental development. *IEEE Transactions on Evolutionary Computation*, 11(6), 2007.

[32] P.-Y. Oudeyer, F. Kaplan, V. V. Hafner, and A. Whyte. The playground experiment: Task-independent development of a curious robot. In D. Bank and L. Meeden, editors, *Proc. of the AAAI Spring Symposium on Developmental Robotics, 2005*, pages 42–47. AAAI, 2005.

[33] M. Prokopenko, F. Boschetti, and A. J. Ryan. An information-theoretic primer on complexity, self-organization, and emergence. *Complexity*, 15(1):11–28, 2009.

[34] M. Rosenblatt. Remarks on some nonparametric estimates of a density function. *Ann. Math. Statist.*, 27(3):832–837, 09 1956.

[35] C. Salge, C. Glackin, and D. Polani. Empowerment–an introduction. In M. Prokopenko, editor, *Guided Self-Organization: Inception*, volume 9 of *Emergence, Complexity and Computation*, pages 67–114. Springer Berlin Heidelberg, 2014.

[36] T. Sauer, J. A. Yorke, and M. Casdagli. Embedology. *J. Stat. Phys.*, 65:579–616, 1991.

[37] H. Sayama. Guiding designs of self-organizing swarms: Interactive and automated approaches. In M. Prokopenko, editor, *Guided Self-Organization: Inception*, volume 9 of *Emergence, Complexity and Computation*, pages 365–387. Springer Berlin Heidelberg, 2014.

[38] J. Schmidhuber. Curious model-building control systems. In *In Proc. Intl. Joint Conf. on Neural Networks, Singapore*, pages 1458–1463. IEEE, 1991.

[39] J. Schmidhuber. Exploring the predictable. In A. Ghosh and S. Tsuitsui, editors, *Advances in Evolutionary Computing*, pages 579–612. Springer, 2002.

[40] N. M. Schmidt, M. Hoffmann, K. Nakajima, and R. Pfeifer. Bootstrapping perception using information theory: Case studies in a quadruped robot running on different grounds. *Advances in Complex Systems*, 16(02n03):1250078, 2013.

[41] R. Shaw. *The dripping faucet as a model chaotic system*. Aerial Press, Santa Cruz, 1984.

[42] J. E. Shore and R. W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Transactions on Information Theory*, IT-26(1):26–37, corrections in IT–29 (6) (1983), 1980.

[43] Y. Sinai. Kolmogorov-sinai entropy. *Scholarpedia*, 4(3):2034, 2009.

[44] F. Takens and E. Verbitski. Generalized entropies: Rényi and correlation integral approach. *Nonlinearity*, 11(4):771, 1998.

[45] J. P. van der Weele and E. J. Banning. Mode interaction in horses, tea, and other nonlinear oscillators: The universal role of symmetry. *American Journal of Physics*, 69(9):953–965, 2001.

[46] O. Vasicek. A test for normality based on sample entropy. *Journal of the Royal Statistical Society. Series B (Methodological)*, 38:54–59, 1976.

[47] X. R. Wang, J. M. Miller, J. T. Lizier, M. Prokopenko, and L. F. Rossi. Quantifying and tracing information cascades in swarms. *PLoS ONE*, 7(7):e40084, 07 2012.

[48] K. Zahedi, N. Ay, and R. Der. Higher coordination with less control – A result of information maximization in the sensorimotor loop. *Adaptive Behavior*, 18(3-4):338–355, 2010.

A Comparison and validation of estimators using autoregressive model of the Lorenz attractor

In this appendix we compare the estimates based on the correlation integral with the estimates using the KSG algorithms for a system for which the excess entropy can be computed analytically: an autoregressive (AR) model of the Lorenz attractor, which results in a stochastic system with defined correlation structure. The dimension of a stochastic systems is infinite, i. e. for embedding m the correlation integral yields the full embedding dimension. However, the excess entropy stays finite.

The AR model of a time series x_i is given by

$$x_{n+1} = \sum_{k=0}^{p} a_k x_{n-k} + \xi_n \tag{42}$$

where p is the order and the a_k are the parameters which are fit to the data and ξ_n independent identically distributed Gaussian noise.

Let us consider the example of an AR2 (p=2) process

$$x_{n+1} = a_1 x_n + a_2 x_{n-1} + \xi_n \tag{43}$$

and calculation the mutual information of subsequent time steps:

$$I\left(X_{n+1} : X_n\right) = 2H\left(X_n\right) - H\left(X_{n+1}X_n\right) \tag{44}$$

$$H\left(X_n\right) = \frac{1}{2}\ln\left(2\pi e\left\langle x_n^2\right\rangle\right) \tag{45}$$

$$\left\langle x^2\right\rangle = a_1^2\left\langle x^2\right\rangle + a_2^2\left\langle x^2\right\rangle + \left\langle \xi^2\right\rangle \tag{46}$$

$$\beta := \left\langle x^2\right\rangle = \sigma^2/\left(1 - a_1^2 - a_2^2\right) \qquad \text{with } \sigma^2 := \left\langle \xi^2\right\rangle \tag{47}$$

$$H\left(X_n\right) = \frac{1}{2}\ln(2\pi e\beta) \tag{48}$$

$$H\left(X_{n+1}X_n\right) = \frac{1}{2}\ln\left((2\pi e)^2|K|\right) \tag{49}$$

$$K = \begin{pmatrix} \beta & \langle x_{n+1}x_n\rangle \\ \langle x_{n+1}x_n\rangle & \beta \end{pmatrix} \tag{50}$$

$$\langle x_{n+1}x_n\rangle = a_1\left\langle x_n^2\right\rangle + a_2\left\langle x_{n-1}x_n\right\rangle \tag{51}$$

$$= \frac{a_1}{(1-a_2)}\beta = r\beta \qquad \text{with } r := a_1/\left(1-a_2\right) \tag{52}$$

$$H\left(X_{n+1}X_n\right) = \frac{1}{2}\ln\left((2\pi e)^2\beta^2\left(1-r^2\right)\right) \tag{53}$$

$$I\left(X_{n+1} : X_n\right) = \frac{1}{2}\ln\left(1-r^2\right) = \frac{1}{2}\ln\left(1 - \left(\frac{a_1}{1-a_2}\right)^2\right). \tag{54}$$

For larger time horizons it suffices to calculate $I\left(X_{n+1}X_n : X_{n-1}X_{n-2}\right)$ because of the Markov property more terms will not increase I.

$$I\left(X_{n+1}X_n : X_{n-1}X_{n-2}\right) = 2H\left(X_{n+1}X_n\right) - H\left(X_{n+1}X_nX_{n-1}X_{n-2}\right) \tag{55}$$

$$H\left(X_{n+1}X_nX_{n-1}X_{n-2}\right) = \frac{1}{2}\ln\left((2\pi e)^4|K|\right) \tag{56}$$

$$K = \begin{pmatrix} \beta & r_1\beta & r_2\beta & r_3\beta \\ r_1\beta & \beta & r_1\beta & r_2\beta \\ r_2\beta & r_1\beta & \beta & r_1\beta \\ r_3\beta & r_2\beta & r_1\beta & \beta \end{pmatrix}. \tag{57}$$

Where the r_i are the correlation coefficient for i-step delay ($r_1 = r$).

$$\langle x_{n+1}x_{n-1}\rangle = a_1\langle x_nx_{n-1}\rangle + a_2\left\langle x_{n-1}^2\right\rangle = a_1r\beta + a_2\beta = \beta r_2 \tag{58}$$

$$r_2 = a_1r + a_2 \tag{59}$$

$$\langle x_{n+1}x_{n-2}\rangle = a_1\langle x_nx_{n-2}\rangle + a_2\langle x_{n-1}x_{n-2}\rangle = a_1r_2\beta + a_2r\beta = \beta r_3 \tag{60}$$

$$r_3 = a_1r_2 + a_2r\,. \tag{61}$$

Thus the 2-step mutual information is given by:

$$I\left(X_{n+1}X_n : X_{n-1}X_{n-2}\right) = \frac{1}{2}\ln\frac{1}{(1+a_2)^2\left((a_2-1)^2-a_1^2\right)}\,. \tag{62}$$

If the process is observed with delay τ, i. e. $y_n = x_{\tau\cdot n}$ then only the r_i in Eq (57) have to be substituted by $r_{\tau\cdot i}$ because β cancels. For convenience the list of the first r_i follows:

$$\begin{aligned}
r_1 &= a_1/(1-a_2) \qquad\qquad r_2 = a_1^2/(1-a_2)+a_2\\
r_3 &= (a_1(-a_1^2+(-2+a_2)a_2))/(-1+a_2)\\
r_4 &= (-a_1^4+a_1^2(-3+a_2)a_2+(-1+a_2)a_2^2)/(-1+a_2)\\
r_6 &= (-a_1^6+a_1^4(-5+a_2)a_2+3a_1^2(-2+a_2)a_2^2+(-1+a_2)a_2^3)/(-1+a_2).
\end{aligned}$$

For larger delays values the expressions become very large such that it becomes intractable to compute the MI for $\tau > 10$ (requires r_{30}) because numerical errors accumulate even after algebraic simplification. The supplementary (playfulmachines.com/QuantBeh2015) contains a Mathematica file for calculating the MI's up to $\tau = 10$.

The predictive information (18) can be also written as

$$\mathrm{PI}_m = \sum_{k=0}^{m-1} 2h_k - \sum_{k=0}^{2m-1} h_k\,. \tag{63}$$

For the AR2 process we have $h_k = h_2$ for $k \geq 2$ due to Markov order 2. Thus, $\mathrm{PI}_1 = h_0 - h_1 = 2H_1 - H_2$ and $\mathrm{PI}_2 = \mathrm{PI}_k = h_0 + h_1 - 2h_2 = 3H_2 - 2H_3$ for $k \geq 2$. For the excess entropy Eq (15) we find $E_1 = 0$, $E_2 = h_0 - h_1 = \mathrm{PI}_1$, $E_3 = E_k = h_0 + h_1 - 2h_2 = \mathrm{PI}_2$, for $k \geq 3$. Fig 9 shows the results for the AR2 model obtained from fitting the parameters to the lorenz system ($a_1 = 1.991843, a_2 = -0.994793$). Indeed the conditional entropies $h_m = h_2$ become identical for $m \geq 2$, see (c),(d) and the excess entropies converge for $m \geq 3$ to their theoretical values, see (e),(f) In the case of the KSG algorithm the values are correct for $m = 1, 2$ but for higher embedding dimensions we get a significant underestimation, see (g),(h). This effect is even stronger for smaller neighborhood sizes k.

B Details of the Decomposition Algorithm

B.1 Fitting

In deterministic scaling ranges the $\delta h_m(\varepsilon)$ follow the form $o - s\log(\varepsilon)$ (possibly with s=0). In a perfectly stochastic scaling range we have $\delta h_m(\varepsilon) \propto$ const for all m. However, the behavior of a system may show several deterministic and stochastic scaling ranges. So we have to determine segments where a certain typical behavior is found. The following algorithm assumes that for each m the $\delta h_m(\varepsilon)$ contain segments of the form $f(\varepsilon) \approx 0 - s\log(\varepsilon)$ and some parts that do not fit. Because for different m and different systems the curves may have different amount of noise we first determine a quality threshold for the fits. Internally the data is represented by n points which we denote $(\varepsilon^i, \delta h^i)$, with $i = 1, \ldots, n$. We denote the range where a fit was computed by $r = (i_l, i_u)$ where $i_l < i_u$. The quality measure is the sum of the squared fit-residuals:

$$q(r) = \sum_{i\in r}(\delta h(\varepsilon^i) - f(\varepsilon^i))^2\,. \tag{64}$$

First, all fits with a certain number of data points (here 10) are computed and their quality measure is computed: $Q = \{q(r_j) \mid |r_j| = 10\}$, where $|r| = i_u - i_l$ is the length of the segment in number of points. There will be segments with low residual errors and segments with very high residuals, namely those that cover the regions between ideal-typical behavior. The heuristics for the quality threshold is the 25% percentile of Q plus a small portion of the standard deviation: $q_{\max} = P(Q, 0.25) + 0.1\mathrm{StdDev(Q)}$. The reasoning for the percentile is to pick a threshold among the good fits, and we only assume that at least 25% of the curve has fitting segments. If there is however an extended range of good fits then a high percentage (say 90%) of short segments have very similar low q values. In this case a too low threshold would be selected which cuts away good fitting regions. The tens of the standard deviation is added to lift the threshold above the potential plateau, see Fig 10(b).

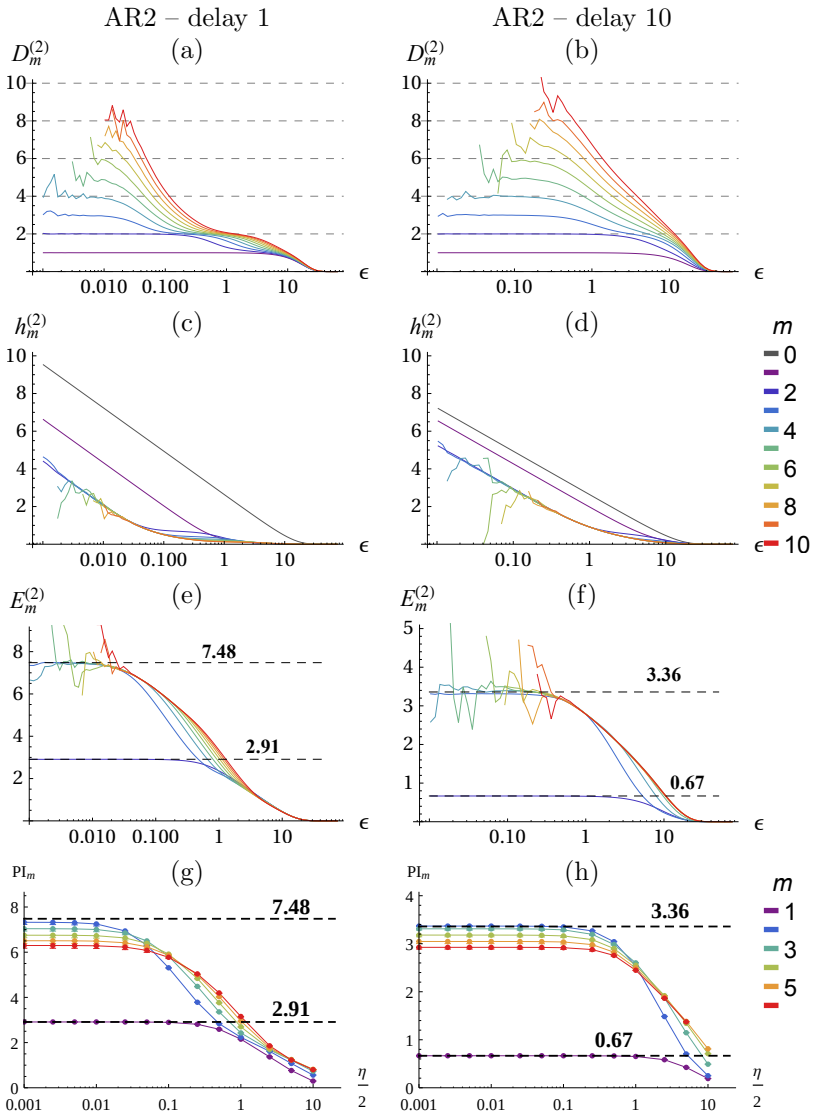

Figure 9: Analysis of an AR_2-model ($a_1 = 1.991843, a_2 = -0.994793$) of the Lorenz attractor observed at delay 1 and delay 10. See Fig 2 for details. **(a)**, **(b)** The dimensionality corresponds to the embedding dimension (as long as the data suffices). **(c)**, **(d)** Conditional entropies show the entropy of the noise for $m > 2$. The difference between first and second, and second and third embedding reveal the actual structure of the time-series, better visible in (e) and (f). **(e)**, **(f)** The excess entropy converges to the theoretical values of 2.91 for $m = 1$ and to 7.48 for $m \geq 2$ (delay 1) and 0.67 for $m = 1$ and to 3.36 for $m \geq 2$ (delay 10). **(g)**, **(h)** Predictive information estimated with KSG algorithm with $k = 20$.

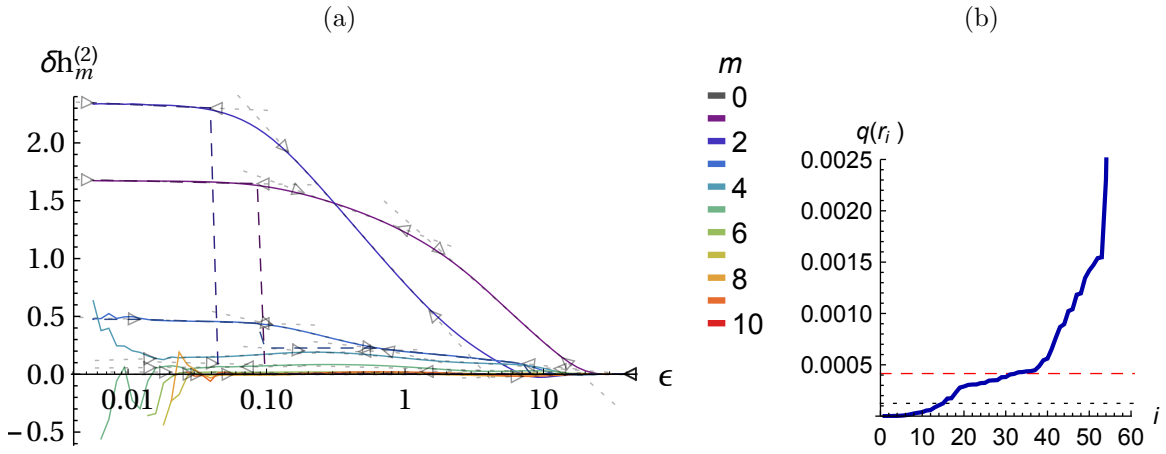

Figure 10: Illustration of the fitting algorithm and the determination of c_m^{MT} and $q_{\max}$. **(a)** Shown are the δh_m curves for the Lorenz system with 0.005 dynamic noise, see Fig 2. The fits are shown by dotted gray lines and their validity range r is marked by triangles. The dashed colored lines show c_m^{MT} for $m = 1, 2, 3$. **(b)** Ranked fit qualities for all segments with 10 data points of δh_1 (solid line) with 25% percentile (dotted line) and threshold $q_{\max}$ (dashed red line). For visibility y-axis is cut, $\max q = 0.013$.

For all "good" segments $Q_{\mathrm{good}} = \{q(r_j) < q_{\max} \mid q(r_j) \in Q\}$ the longest extension of the fitting range is determined that keeps the quality of the fit below the threshold, i. e. $\forall q(r_j) \in Q_{\mathrm{good}}, (i_l, i_u) = r_j : i'_u = \arg\max_{k>i_u} q((i_l, k)) < q_{\max}$. Thus we get many overlapping regions of good fits. For each pair of two regions that overlap more than 30% (w.r.t. the smaller) we discard the shorter one. The remaining regions are our final fits, as displayed in Fig 10(a).

The components c^{MT} for calculating the constant of the middle term (37) are also displayed in Fig 10(a). For $m = 3$, the value of c^{MT} is 0.23 for $0.1 < \varepsilon < 0.8$ because of the shallow region at $[0.8, 9]$. This leads to the value of 0.68 ($m * c_m^{\mathrm{MT}}$) for the state complexity in Fig 3(b). For $\varepsilon < 0.1$ the constant c_3^{MT} takes the value of ≈ 0.5 of the next (smaller scale) plateau. Similarly $c_{1,2}^{\mathrm{MT}}$ raise from zero to the value of the plateaus, which leads to the increase in state-complexity in the stochastic scaling range. Remember that in the stochastic scaling range ($\varepsilon < 0.1$) $m_l = 1$ and $m_u = 3$ are taken from the preceding deterministic scaling range.